

Selecting Pixels for Kepler Downlink

Stephen T. Bryson^a, Jon M. Jenkins^b, Todd C. Klaus^c, Miles T. Cote^a, Elisa V. Quintana^b,
Jennifer R. Hall^c, Khadeejah Ibrahim^c, Hema Chandrasekaran^{a,b}, Douglas A. Caldwell^b,
Jeffrey E. Van Cleve^b, Michael R. Haas^a,

^aNASA Ames Research Center, M/S 244-30, Moffett Field, CA USA 94305

^bSETI Institute/NASA Ames Research Center, M/S 244-30, Moffett Field, CA USA 94305

^cOrbital Sciences Corporation/NASA Ames Research Center, M/S 244-30, Moffett Field, CA
USA 94305

ABSTRACT

The *Kepler* mission monitors $> 100,000$ stellar targets using $42\ 2200 \times 1024$ pixel CCDs. Bandwidth constraints prevent the downlink of all 96 million pixels per 30-minute cadence, so the *Kepler* spacecraft downlinks a specified collection of pixels for each target. These pixels are selected by considering the object brightness, background and the signal-to-noise of each pixel, and are optimized to maximize the signal-to-noise ratio of the target. This paper describes pixel selection, creation of spacecraft apertures that efficiently capture selected pixels, and aperture assignment to a target. Diagnostic apertures, short-cadence targets and custom specified shapes are discussed.

Keywords: *Kepler*, exoplanet, transit, pixel selection

1. INTRODUCTION

The *Kepler Mission* continuously observes $\sim 165,000$ target stars in *Kepler's* 115 square degree Field of View (FOV) seeking to discover Earth-like planets transiting solar-like stars by detecting photometric signatures of transits.^{1,2} Data is collected and stored for monthly downlink, and the data is processed in the Science Operations Center (SOC).^{3,4}

Every target is observed with CCD readout every 6.52 seconds and co-added into 29.4 minute observations, referred to as *Long Cadence* (LC) data. A smaller number of targets, at most 512, is co-added into 58.8 second *Short Cadence* (SC) observations. These data are collected nearly continuously for about 30 days and downlinked via high-bandwidth Ka-band transmissions. Bandwidth constraints of the Ka-band transmission limits the amount data that may be downlinked, making it impossible to downlink all 96 million pixel values that are collected with each LC observation. It is therefore required to identify the pixels that provide the best data for the target stars and select these pixels for downlink. This paper describes the method by which the required pixels for each target are determined.

Kepler has been placed in an Earth-following Solar orbit. Every ~ 90 days the spacecraft rotates 90 degrees, referred to as a *quarterly roll*, which changes which stars are on which CCDs. Every quarter the pixel selection process is repeated with a new collection of targets.⁴

One of the primary goals of pixel selection is to identify pixels that are optimal for aperture photometry, as described in §2. The high quality of the *Kepler* photometric data⁵ indicates that such pixel selection has been basically successful.

Further author information: Send correspondence to S.T.B.: E-mail: Steve.Bryson@nasa.gov

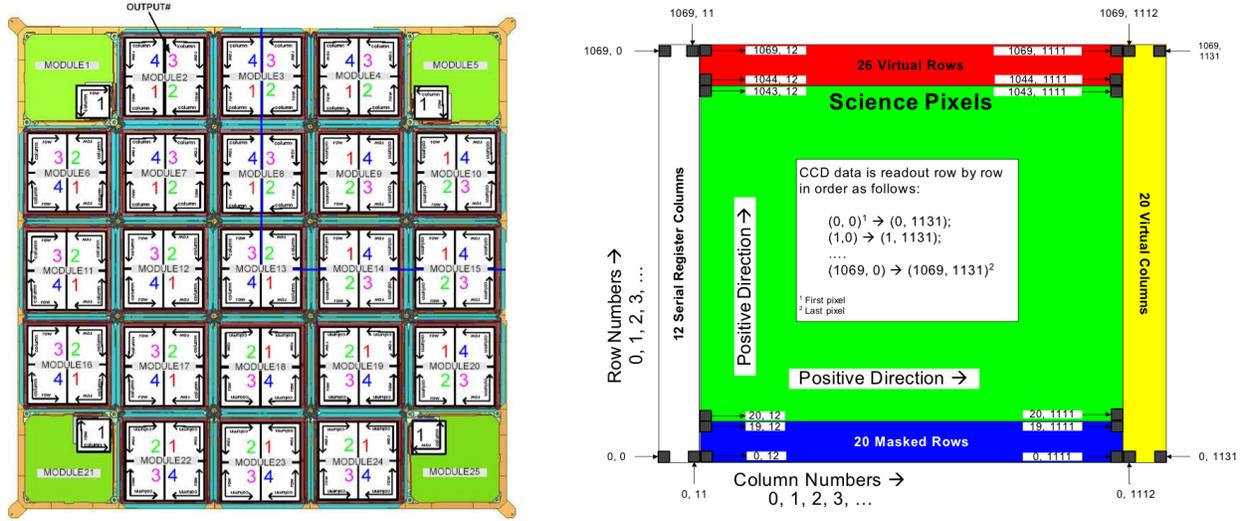


Figure 1. Left: The CCD array on the *Kepler* focal plane, showing the 21 modules, each of which has 2 2200 × 2048 pixel CCDs. Each CCD is read out via 2 output channels. Smaller CCDs used as fine guidance sensors are also shown in each corner, but are not discussed in this paper. Right: the pixel arrangement of each output channel.

1.1 The *Kepler* focal plane

The *Kepler* focal plane science sensors consist of 42 2200 column by 1044 row CCDs mounted on 21 electronic modules (Fig 1) with an image scale of 3.98" per pixel. The first 20 rows of each CCD are not exposed to the sky in order to provide calibration and diagnostic data as described below. Each CCD is divided into two 1112 × 1044 *output channels*. Each output channels is supplemented by 26 trailing virtual rows, 12 leading serial register columns and 20 trailing virtual columns, giving each output channel 1136 × 1070 addressable pixels. The 12 leading serial register columns and 20 trailing virtual columns are used to collect black level data for each row. *Kepler's* lack of a shutter means that pixels are exposed to the sky during readout, which causes image smear along columns. The leading 20 masked and 26 trailing virtual columns measure this smear data. The black level and smear data are called *collateral data* and are used to calibrate the pixel data during ground processing.⁶ The pair of CCDs on each module provide a contiguous 2200 × 2048 pixel image of a portion of the *Kepler* field of view.

1.2 Target and aperture definition task flow

Kepler on-board flight system software does not allow specification of arbitrary collections of pixels for each target. A table of 1024 *aperture masks* is loaded onto the spacecraft, and the pixels associated with a target are fit into one of these masks. Efficient design of these masks to fit the required pixels without too many excess pixels is a primary concern of the work described in this paper.

The pixel selection workflow proceeds along the following lines.

1. For each target required pixels are selected (§2).
2. A mask table is created to match the set of selected pixels for all targets (§3.2).
3. Each target's required pixels is assigned to the smallest mask that contains all required pixels (§3.3).

One target, particularly very bright stellar targets, may be assigned more than one mask.

The resulting mask assignment for each target is collected into *target definitions* that include the index in the aperture mask table of the selected aperture mask and the location on the focal plane (module, output, row, column) of that mask.

1.3 *Kepler* pixel and target types

Kepler pixels are collected for several types of targets:

Stellar targets are point-like sources⁷ whose pixels are selected to maximize the signal to noise ratio (SNR) (§2.2). Stellar targets are specified by a Kepler ID, which is used to look up pertinent data in the Kepler Input Catalog (KIC).⁸ Stellar targets may be either LC or SC.

Custom targets are explicitly specified collections of pixels. Custom targets are defined by a reference pixel position and a set of offsets, one for each pixel, from that reference position. Custom targets are used for non-stellar sources and diagnostic collections of pixels, and may be either SC or LC.

Background targets are small (nominally 2×2) sets of pixels that sample the background signal in long cadence. These pixels are selected to support a 2D polynomial representation of the background (§2.3).

Reference pixel (RP) targets are special stellar targets used for diagnostics whose pixels are downlinked bi-weekly via low-bandwidth X-band communications.⁹ Reference pixel targets are designed to include sufficient information for calibration and analysis (§4), so special treatment of the

LC, SC, background and RP targets are treated separately by the flight system software, with their own memory allocation. LC, SC and RP share the same mask table, and SC and RP targets also appear on the LC target list. Background targets have a separate mask table.

1.4 Pixel selection requirements

Memory, bandwidth and flight software design impose several constraints on the final set of pixels selected for downlink:

Long cadence: There may be no more than 170,000 target definitions, with no more than 10,000 per output channel, and no more than 5.44 million pixels. This implies an average of ≤ 32 pixels per target definition to attain 170,000 target definitions.

Short cadence: There may be no more than 512 target definitions, and no more than 43,520 pixels or an average of ≤ 85 pixels per target if 512 targets is desired.

Background: There may be no more than 1125 target definitions and no more than 45,000 pixels (for an average of 4 pixels per target) for each of 84 output channels.

Reference pixels: At the beginning of the *Kepler* mission, there may be no more than 96,000 pixels in a set of reference pixel targets across the entire focal plane. As *Kepler* moves away from Earth bandwidth degradation reduces the number of reference pixels that can be downlinked in a bi-weekly contact.

Aperture mask table: Up to 1024 aperture masks may be defined, using no more than 87,040 pixels. This implies an average of 85 pixels per mask for 1024 masks.

For reasons described in §4, 252 entries in the aperture mask table are dedicated to supporting RP observations, leaving 772 aperture masks for LC and SC targets.

Basic *Kepler* mission requirements specify the capability to observe stars with magnitudes between 9 and 15. Simultaneously meeting the above constraints yet providing efficient mask assignment for a wide range of stellar magnitudes is a significant challenge. Observation of stars brighter than magnitude 9 is allowed with the understanding that pixel efficiency will be significantly degraded, so these observations can use a very large number of pixels.

2. PIXEL SELECTION

Pixels for stellar targets are selected to maximize the SNR for that target, while background pixels are selected to facilitate the construction of 2D polynomial representation of the background. Both pixel types are selected based on a model of the signal seen by *Kepler* CCDs created from observed characteristics of the sky and *Kepler* photometer including optics and electronic properties. This model is used to create a synthetic sky image, which is analyzed to determine the desired pixels. The basic strategy for optimal pixel selection is to create two synthetic images: one with all stars in the region of the target star, and another with all stars except the target star itself. These images are then used to compare the signal from the target star with the noise from the target star, stellar background, and the instrument.

2.1 Synthetic Image Creation

The synthetic image is created using the following elements:

Kepler Input catalog providing J2000 right ascension, declination, and magnitude in the *Kepler* bandpass of stellar targets in the *Kepler* field.

Pixel Response Function (PRF)¹⁰ model, an observation-based super-resolution model of how light from a star falls on *Kepler* pixels at different locations in the focal plane. There is one PRF model for each output channel, and this model contains 5 PRFs that are linearly interpolated to capture intra-channel PRF variations. The observations behind the PRF model were taken at a 15-minute cadence, sufficient to capture the LC behavior of spacecraft pointing jitter. The PRF model includes intra-pixel variability, intra-channel variation of the PRF

Focal plane geometry (FPG) and pointing model,¹¹ which includes measurements of the locations of the CCDs in the *Kepler* focal plane, models of the Kepler optics and of differential velocity aberration (DVA). These models are used to determine the pixel location of the central ray of each stellar target. DVA can move a star as much as 0.6 pixels in a quarter.

Saturation model,¹¹ which includes information about the well depth of each output channel.

Zodiacal light model represented as a mesh of magnitude values on the sky.

Read noise model,¹¹ observed values of read noise for each output channel.

Charge transfer efficiency CTE model, which describes how much flux is lost with each charge transfer during readout.

The synthetic image used for pixel selection models the signal in calibrated pixels, so smear due to the lack of a shutter and other instrumental effects are not included. The pixels selected for optimal photometry in §2.2 are used for an entire quarter of about ~ 90 days, so for every star the the pointing model is used to smear the PRF along the path taken by that star due to DVA during that quarter.

The synthetic image for each output channel is generated star by star. For each star in the KIC that falls on the output channel,

1. The star's pixel position on the channel is computed from the star's RA and Dec in the KIC, using the FPG and pointing model (including sub-pixel position).
2. The PRF at the star's position is evaluated along the path taken by the star due to DVA, creating a smeared PRF covering the entire quarter.
3. The pixels resulting from the star's PRF are normalized to sum the flux of the star derived from the *Kepler* magnitude in the KIC. The resulting *target-only image* is saved for each target for use in §2.2.
4. The normalized pixels are added to the synthetic image at the appropriate pixel location.

At this point the synthetic image represents the collection of stars in the KIC as they would appear on the CCD without any saturation or zodiacal light background. At this point a copy of the synthetic image, called the *background image* is made for use in the optimal pixel selection described in §2.2. The zodiacal light is interpolated onto each pixel and added to the synthetic image. Saturation is then iteratively spilled along columns by moving flux exceeding the well depth up and down the column, with the fraction of flux moving up vs. down being controlled by an input parameter. Finally the CTE model is applied. An comparison of the resulting synthetic image and the in-flight image of the same pixels is shown in Fig. 2.

There are several sources of error in the generation of the synthetic image, which can result in compromised photometrically optimal apertures:

PRF errors: Each PRF model is computed as an average of several observed stars without consideration of color. The actual PRF of individual stars will differ because the optical PSF underlying the PRF has minor color dependence. The PRF varies within a channel, with stronger variation near the edge of the FOV, and this variation is only approximately modeled by the linear interpolation included in the PRF model.

KIC errors: The KIC contains several artifacts which do not correspond to objects in the real sky. Such artifacts can be seen in Fig. 2 around the bright star at (520, 580), where there is a faint diagonal line of stars extending from the lower left to the upper right through the core of the star. This is likely a diffraction spike mis-identified as faint stars and does not appear in the flight image. The KIC does not reflect stellar variability, which is not known for many stars in the *Kepler* field (for stars whose variability is known the brightest magnitude is used). An example of such magnitude errors can be seen in the star at (540, 570). When a KIC underestimate of the a target’s magnitude has been identified, there is a mechanism to provide corrected magnitudes to the pixel selection process. of Fig. 2 which is brighter in the synthetic image than in the flight image.

Saturation model: The saturation model in use at this time of writing is very simple, and provides the well depth per output channel, assuming that the well depth does not vary within the channel and that the spill up and down a column is symmetric. Both of these assumptions are strongly violated, with variations in both well depth and symmetry. These variations have a spatial frequency higher than our ability to measure them with the bright stars in the *Kepler* FOV. To provide margin against this uncertainty in saturation spill asymmetry the simulated saturation is extended by 50% in both directions. An approach of explicitly providing the saturation spill per bright star is being investigated.

Changing focus: The focus of the *Kepler* photometer has been observed to undergo seasonal changes as the sun angle varies throughout the *Kepler* orbit. In addition smaller, high-frequency focus changes have been observed that are highly correlated with various heaters operating on the spacecraft. These focus changes cause PRF changes not captured in the synthetic image at this time. Focus changes also induce plate scale variations, so the stars will not be placed in the correct positions in the synthetic images. Though plate scale-induced position errors are small their effects have been observed in photometry and corrective measures are being implemented in the SOC computation of the synthetic images

Because these errors were anticipated, the mask assignment process places a 1-pixel halo around the optimal aperture described in §2.2 prior to aperture mask fitting, as described in §3.

2.2 Optimal pixel selection

This target-only image created in §2.1 is subtracted from the background image, creating the background image with all stars on this channel except the target. Saturation and CTE are then simulated as described in §2.1. Pixel values p_{target} in the target-only image define the signal of each pixel, while the pixel values p_{back} in the background image provides the background signal. The SNR of each pixel is estimated as

$$SNR_{\text{pixel}} = \frac{p_{\text{target}}}{\sqrt{p_{\text{target}} + p_{\text{back}} + \nu_{\text{read}}^2 + \nu_{\text{quant}}^2}} \quad (1)$$

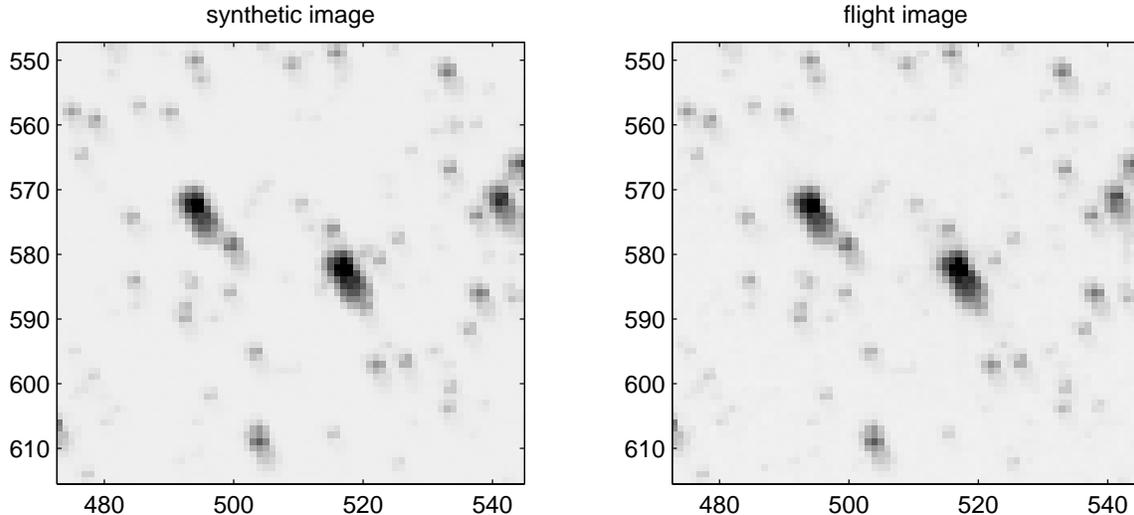


Figure 2. A comparison of a synthetic image (left) with the actual flight image of the same region of the sky, near the edge of the *Kepler* where the PRF is large. For purposes of comparison with the sky image the synthetic image is computed for a short time interval for the time of the flight image so the PRFs are not smeared.

where ν_{read} is the value for this channel from the read noise and ν_{quant} is the quantization noise. Quantization noise is given by

$$\nu_{\text{quant}} = \sqrt{\frac{n_C}{12} \left(\frac{w}{2^{n_b-1}} \right)^2} \quad (2)$$

where n_C is the number of cadences in a co-added observation (30 for LC), w is the well depth and n_b is the number of bits in the analog-to-digital converter ($= 14$). This noise formula (1) includes Poisson shot noise of both the target and background pixel values.

Given a collection of pixels, the SNR of the collection is given by the RSS of the SNR of the component pixels. Optimal pixel selection begins by including the pixel with the highest SNR. The next pixel to be added is the pixel that results in the greatest increase in SNR of the collection. Initially the collection SNR will increase as pixels are added. After the bright pixels in the target have been added, dim pixels dominated by noise cause the collection SNR to decrease. The pixels collection with the highest SNR defines the optimal aperture. Fig. 3 shows an example of the dependence of SNR on pixel inclusion.

The background and target images allows us to estimate crowding by estimating the fraction of flux in the optimal aperture that is due to the target star. The resulting *crowding metric* is useful for estimating the dilution of flux from the target in the optimal aperture, which has an impact on the detectability of transits.⁷ The same measure on a 21×21 pixel square provides a *sky crowding metric* which is useful for identifying uncrowded stars. The fraction of target's flux that falls in the optimal aperture is also computed.

2.3 Background pixel selection

Background targets are created in a way that best supports the creation of a 2D polynomial representation of the background on each channel. This background polynomial is subtracted from the pixel values prior to aperture photometry.¹² To lead to good polynomials the background targets should be homogeneously and uniformly distributed on the output channel. Because the accuracy of background polynomials increases with an increasing density of data points but diminishes near the edge of the polynomial domain, more background targets are placed at the edges of the output channel. To achieve this distribution the intersections of an irregular Cartesian mesh are used for the initial guess at background target positions. Because this mesh is the product of two linear meshes we cannot get a homogeneous distribution and exactly 1125 targets, but a choice of $31 \times 36 = 1116$ mesh lines is close.

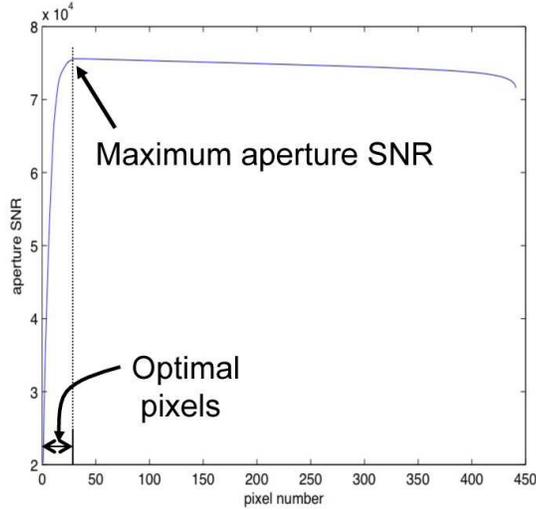


Figure 3. The aperture SNR curve built up as pixels are added in order of decreasing pixel SNR. The optimal set of pixels is defined as the set of pixels required to reach the maximum of this curve.

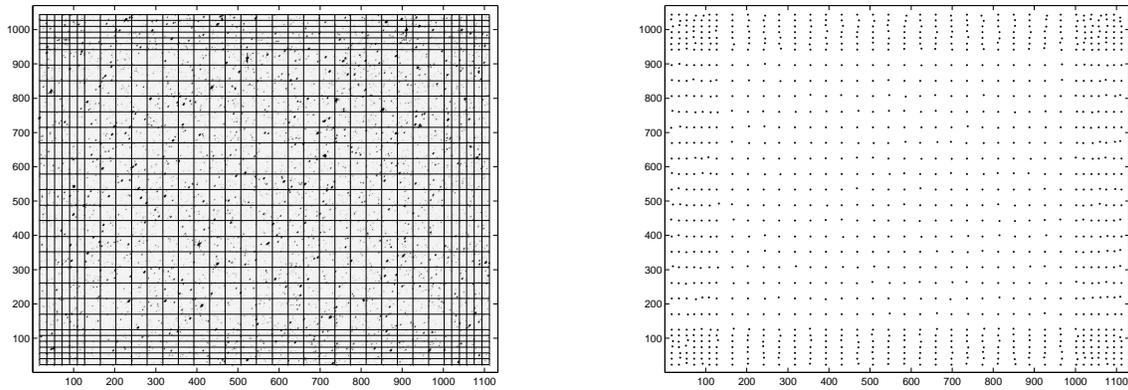


Figure 4. Left: The grid used to seed background target positions on an output channel. Right: Final locations of 2×2 pixel background targets.

A pixel in the synthetic image is considered to be background if its value is less than the dominant mode of the pixel histogram, and a 2×2 pixel background target is considered valid if all four pixels are below the background threshold. For each mesh intersection a background target is placed. If that target is not a valid background target, increasingly larger boxes centered on the mesh intersection are searched until either a valid background target location is found or the box exceeds a maximum size. In the latter case the target with the smallest summed pixel value is chosen. The performance of this approach is enhanced by creation of a 2×2 -average binned image, so the search in each box is done by taking the minimum of the binned image in that box. The initial mesh and resulting background locations are shown in Fig. 4.

3. MASK CREATION AND ASSIGNMENT

The pixel selection process in §2 combined with custom apertures produces several thousand uniquely shaped pixel apertures across the *Kepler* field of view. These shapes must be mapped to 772 aperture masks loaded onto the spacecraft in a way that minimizes the number of excess pixels that are collected. This presents a difficult combinatorial problem, whose optimal solution was not supported by available resources during Kepler

development. A simpler, near-optimal approach was developed that used performs an iterative statistical estimate of which pixel apertures may be best used as entries in the mask table.

The approach used produces aperture masks that are efficient for dimmer targets where the vast majority of pixels are found, but do a relatively poor job for bright targets. This is addressed by creating several manually-chosen masks for bright targets.

The mask creation and assignment process proceeds in three stages: final required pixel apertures are defined by adding pixel margin to the requested pixels. The mask table is created from analysis of the resulting final set of pixel apertures, supplemented by manual mask creation for bright targets. Finally the required pixels per target are mapped to aperture masks in the mask table, creating the final target definitions. The mask table and target definitions are then delivered for quarterly upload to the spacecraft.

3.1 Required pixels: adding pixel margin

Once the pixels for a target have been specified, either by the optimal pixel selection process in §2 for stellar targets or by explicit specification for custom targets, the target may be assigned extra pixels. These extra pixels are of two types:

Halos are rings of pixels around the specified pixels for a target, typically applied as margin against uncertainty.

A pixel is added to the halo if any of the 8 adjacent pixels including corners are included in the target’s specified pixels. As many as 4 halos may be added, which are added iteratively so, for example, the third halo treats the previous two halos as specified pixels for the target.

Undershoot column is an extra column of pixels to the left (“upstream” in the pixel readout) of the specified pixels to provide data for the undershoot correction algorithm.⁶

Halos and the undershoot column can be specified on a target-by-target basis, and when both are present the halos are applied first. The default for stellar targets is to add one halo and the undershoot column, and the default for custom targets is no halo or undershoot column.

The final pixels with the specified pixels and any halo or undershoot column pixels are called *required pixels*.

3.2 Mask table creation

We break the table of 1024 aperture masks into 4 sections:

Dim target masks N_{dim} masks that are algorithmically generated by an iterative statistical analysis of the required targets, creating masks that fit required apertures with the smallest number of excess pixels.

Dedicated masks N_{dedicate} masks that automatically set equal to the required pixels for specific targets. These masks are used for oddly-shaped diagnostic targets or bright high-value stars that are difficult to fit efficiently using other masks.

Bright target masks N_{bright} masks that are specially generated to fit long saturated columns or large cores of bright stars.

Reference pixel target set masks $N_{\text{RPTS}} = 252$ masks that are specifically designed to collect collateral data (black level and smear) for reference pixel targets, as described in §4.

The dim target portion of the mask table is the largest section of the mask table (currently $N_{\text{dim}} = 736$) and is filled in using the following algorithm:

1. The aperture mask table is initialized to contain simple geometric shapes.
2. Mask assignment is performed using the full target set as described in §3.3.
3. The unique required apertures are identified.

4. Masks that are perfect fits to some required aperture are identified and set aside.
5. Masks that were not perfect fits are sorted in descending order of the total number of excess pixels that are associated with that mask.
6. The masks in the last step are replaced with shapes from the unique set of required apertures for targets dimmer than a specified magnitude until the dim target region of the mask table is filled.

This process is iterated twice in nominal use. The result is the dim target portion of the mask table containing masks appropriate for dim stars, which make up the bulk of the target set.

The bright target portion of the mask table (currently $N_{\text{bright}} = 28$) is filled in using two methods. First, the magnitude range between 7 and 11 are divided into several magnitude intervals (for the current mask table 18 intervals are used). Within each interval the required pixels of all targets with *Kepler* magnitude within that interval are combined to make a single aperture mask that fits all targets within that interval. For targets brighter than magnitude 7 masks are hand-specified to capture their saturation and large cores.

The dedicated mask portion of the mask table is empty prior to final mask assignment, and is filled in as described in §3.3.

3.3 Mask assignment

Once the mask table is complete they can be assigned to target definitions. Because SC targets are also on the LC target list, only the LC list needs to be considered. RP targets are assigned masks separately as described in §4.

The majority of target definitions are assigned aperture masks by choosing the mask that contains the fewest pixels and completely includes the target's required pixels. Such a mask and the position of the target definition inside the mask is chosen by convolving the target's required pixels (rotated by 180 degrees) with the aperture mask pixels. If elements of the convolution are equal to the number of required pixels then the target's pixels fit in the mask. The position of the mask relative to the target pixels are determined by the centroid of all pixels in the convolution whose values are equal to the number of required pixels.

If these if no aperture mask that contains the required pixels, the required pixel set is divided in two along the shorter axis and an attempt to find an aperture mask for each piece is performed. This division approach is applied iteratively until masks are found that cover all required pixels.

Dedicated mask targets are not algorithmically assigned masks by the above process. If there is not already an aperture mask that exactly fits the target's required pixels a slot in the dedicated mask region of the aperture mask table (see §3.2) is filled in by the required pixels of this target. This process assumes that the dedicated mask portion of the mask table has been sized to accommodate all dedicated mask targets.

Examples of targets and their assigned masks are shown in Fig. 5. In operations the number of unused pixels downlinked due to masks over-fitting targets is 4%, which are mostly from bright stars. Fig. 6 shows the number of mask pixels per target, the cumulative pixel count and the number of excess pixels per target as functions of target magnitude. We see that the excess pixels are mainly in the bright targets while the dimmer targets, which make up the overwhelming bulk of targets, are well fit.

4. REFERENCE PIXEL TARGETS

The primary design consideration for reference pixels is the very limited number of pixels permitted in an X-band transmission due to low bandwidth. The 96,000 available pixels are used to provide attitude determination and diagnostics on 84 output modules.⁹ The requirement to calibrate stellar targets for this purpose implies that black-level and smear data for each target must be included in the 96,000 pixel budget. Late in the mission this budget becomes much smaller as the bandwidth decreases due to the increasing distance of the spacecraft.

The target distribution strategy for reference pixels is to assign up to 5 bright, unsaturated targets to each of the 84 output channels. Channels near the edge of the FOV are preferentially given 5 targets to increase the attitude determination algorithm's sensitivity to roll. At the beginning of the mission, while bandwidth is high,

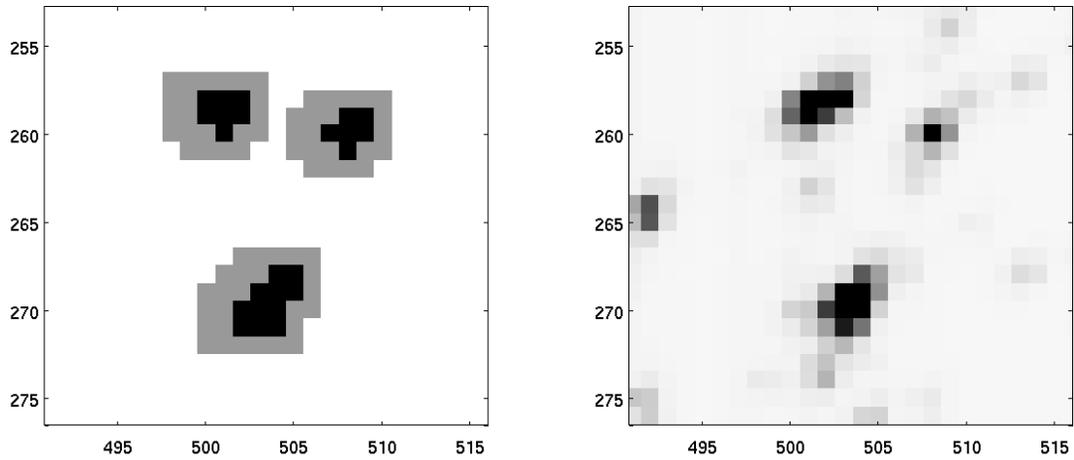


Figure 5. Left: Examples of the masks (grey) selected once the halo and undershoot column is added to the optimal aperture (black). Right: A flight image showing the targets captured by the masks on the left.

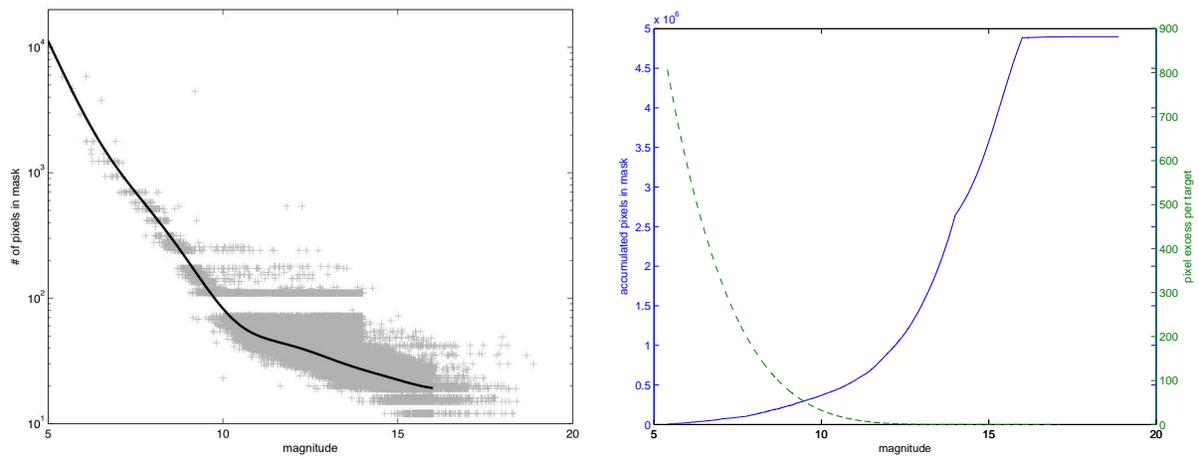


Figure 6. Left: The number of pixels per mask vs. target magnitude for stellar targets including pixel halo and undershoot column. The pixel number for each mask is shown by a grey cross, and a 10-th order robust polynomial fit showing the mode of the pixel number distribution is shown by the black line. The smallest mask (for a 1-pixel optimal aperture) is 4×3 . Banding for the brighter stars is due to the smaller number of masks available for these stars. The smallest mask Right: Cumulative number of pixels in a target's mask vs. magnitude (solid line, left axis), showing that most pixels are on dim targets, and the number of excess pixels per target (dashed line, right axis), showing that bright targets have the largest pixel excess and dim targets have very little pixel excess

all channels have at least 3 reference pixel stellar targets to support plate scale measurements. For the first half of the mission all reference pixel targets are stellar targets. As the bandwidth decreases, smaller rectangular custom targets will be used for diagnostics on some output channels.

For each reference pixel target a few nearby background pixels are chosen, as well as the black and smear values specific to this target. For each output channel a background mask is created that is exactly equal to the the background pixels on the channel. Similarly a smear mask is created to capture the smear pixels on the channel, and a black mask is created to capture the black pixels on each channel. Three masks per channel are required in order to satisfy the requirement that masks contain no more than 85 pixels on average. The stellar target pixels are given two halos and an undershoot column, and are assigned masks using the algorithm described in §3.3.

5. CONCLUSIONS

The algorithms and approaches described in this paper have been successful in specifying the pixels to be downlinked from the *Kepler* spacecraft. These algorithms are used to provide quarterly updates of the desired pixels to the *Kepler* spacecraft. A variety of sometimes conflicting requirements are met, while delivering good performance in terms of good photometric apertures while acquiring a small number of unneeded pixels.

ACKNOWLEDGMENTS

The authors would like to thank the *Kepler* SOC operations team, particularly Brett Stroozas and Jeneen Sommers for their infinite patience and firm guiding hand. The entire SOC development team has been crucial in its support, as has been the entire *Kepler* team. Finally we wish to thank Bill Borucki and David Koch for their leadership of the *Kepler* Mission.

Funding for this Discovery mission is provided by NASA's Science Mission Directorate.

REFERENCES

- [1] Borucki, W., et al., "Kepler planet-detection mission: introduction and first results, *Science*, **327**, pp. 977–980, 2010.
- [2] Koch, D. G. *et al.*, "Kepler mission design, realized photometric performance, and early science," *ApJ Lett.*, **713(2)**, pp. L79–L86, 2010.
- [3] Middour, C. K. *et al.*, "Kepler Science Operations Center Architecture," these proceedings, 2010.
- [4] Haas, M. *et al.*, "Kepler Science Operations," *ApJL* 713 (2010) L115-L119.
- [5] Caldwell, D. A. *et al.*, "Instrument Performance in Kepler's First Months," *ApJL* 713, L92-L96.
- [6] Quintana, E. V. *et al.*, "Calibrations," these proceedings (2010).
- [7] Batalha, N. M. *et al.*, "Selection, Prioritization, and Characteristics of Kepler Target Stars," *The Astrophysical Journal Letters*, **713**, pp. L109–L114, 2010.
- [8] Latham, D. W. *et al.*, "The Kepler input catalog," *Proc. AAS* 207, 1340 (2005).
- [9] Jenkins, J. M. *et al.* "Semi-weekly monitoring of the performance and attitude of Kepler using a sparse set of targets," these proceedings, 2010.
- [10] Bryson, S. T. *et al.* "The Kepler Pixel Response Function," *The Astrophysical Journal Letters*, **713**, pp. L97–L102, 2010.
- [11] Allen, C. A. *et al.*, "Focal Plane Characterization," these proceedings (2010).
- [12] Twicken, J. D. *et al.*, "Photometric Analysis in the Kepler Science Operations Center Pipeline," these proceedings (2010).